Abstract or Die: Life, Artificial Life and (v)organisms

Ron Cottam, Willy Ranson and Roger Vounckx

The Evolutionary Processing Group
Vrije Universiteit Brussel, Belgium

ABSTRACT

One of the questions which most bedevil the development of quasi-autonomous systems is “what is intelligence?” It is fair to say that without an at least provisional definition of intelligence the conception of robotic entities which can satisfactorily mimic our own capabilities will stagnate. But it is far from clear that such a definition can be established other than in a multitude of different contexts. We have established a grounding framework for this multi-contextual definition by linking together the logical extremes of natural phenomena and our *abstract* representation of them. This paper describes the implications of that framework for the creation of virtual (v)organisms, which would be closer relatives to biological organisms than to the “artificially alive” and “artificially intelligent” products of digital computation. We conclude that the primary criterion for creation of a (v)organism is the establishment of a self-correlating multi-leveled linking structure between *natural* and *abstract logics* which is unified through quantum mechanical entanglement. We foresee that, however a (v)organism is initially set up, we will have to progressively develop its ultimate goals through an active co-operation which is closely related to biological parenting.

Keywords: *virtual organism, (v)organism, natural logic, abstract logic, hierarchy, quantum mechanical entanglement.*

1. INTRODUCTION

What is “intelligence”?

James Albus [1] defines it as

> *the ability to act appropriately in an uncertain environment; appropriate action is that which maximizes the probability of success; success is the achievement or maintenance of behavioral goals; behavioral goals are desired states of the environment that a behavior is designed to achieve or maintain.*

David Fogel [2] has suggested that intelligence is

> *the ability of a system to adapt its behaviour to meet its goals in a range of environments*.

We agree that these two descriptions approximate to a “best currently available definition”. The difficulty, however, with these and other “top-down” *natural* language definitions, appears when we try to identify the essentials which will enable us to deconstruct it in the time-honored “reductive” manner and then build an *intelligent* system “bottom-up”. What is “*ability*”? What is “*a system*”? What is “*behavior*”? What *exactly* are “*goals*”? What are “*environments*”? Unfortunately, these descriptions do not readily lend themselves to reductive linguistic processing, as the elements of the complete expressions are inter-dependent: we cannot establish definitions of the words *ability*, *system*, *behaviour*, *goals* and *environments* in isolation and then extract the complete expressions’ meanings by simply combining them.

Bruce Edmonds [3] has presented a related argument as a tentative definition of “complexity”, namely that it is

> *that property of a language expression which makes it difficult to formulate its overall behavior even when given almost complete information about its atomic components and their inter-relations*.

The situation is no better if we resort to *formal* language as our descriptive mode: in fact, it is a good deal worse! Even a humble Boolean AND expression suffers from this irreversibility: we can derive the single output from multiple inputs, but not the complete inputs from the output. Science *itself* is similarly flawed in its relationship to nature, in its use of rationalities which systematically presuppose reversible elemental interchangeability (e.g. as in the presupposition that if A = B + C then B + C = A).

Nature appears to use a different approach, where the possibility of contradiction between representations at different localities is provided by relativity, but even so global coherence is maintained. Local representations which are at odds with the global “picture” are destroyed in favor of their local-globally coherent companions: note the quantum-mechanical collapse of multiply-superimposed hypothetical representations into one “real” conclusion.

We would do well to take account of the difference between our *abstract* formulations, which impose local-global separation and are globally inconsequential, and the fundamentally different formulations of nature, which entrain unavoidable consequences. *We* may be unable to formulate correctly the dynamics of even only three interacting bodies (the Newtonian 3-body problem), but apparently nature suffers from no such difficulty. The present paper examines this difference as part of a route towards building intelligent systems. We will refer to *abstract logic* and *natural logic* to differentiate between the usual (artificial) approaches and those taken by nature.

Digital technology constitutes the *real* implementation of *abstract logic*, where although consequences are intentionally localized they may have disastrous global implications (e.g. system crash). The main target of this paper's considerations is the meaning of the word "goals" which appears in James Albus's and David Fogel's descriptions.

2. WHAT IS A SYSTEM?

A system is by its nature and designation unified. Unification of a dynamic assembly into *a* system is implicit in its representation as a number of nodes or sub-units *which communicate with each other*. The central issue is this *inter-communication*, which explicitly serves to fulfill the (presumed) function of an assembly, but implicitly achieves its integration into a single unit. The selection of an inter-communicational logic defines the degree and effectiveness of system unification.

2.1. Abstract logic is insufficient for unification

The central character of an *abstract logic* is the restriction which is imposed on the globalization of local effects, and *vice versa*. Although *abstract logic* itself is inconsequential, its compartmentalization has global consequences within a logic structure.

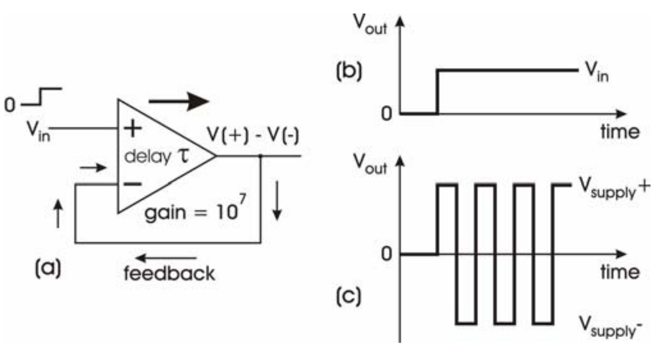

Figure 1: (a) definition of an operational amplifier circuit, (b) the conventional output, and (c) a more likely output.

An excellent example of this is provided by the classic student exercise of determining the output of an operational amplifier which has its output fed back directly to the negative input, as shown in Figure 1(a). Conventional wisdom would suggest that the output of the amplifier defined in Figure 1(a) would be as shown in Figure 1(b) following an input voltage step. Not so. An output which more closely corresponds to the definition of Figure 1(a) is illustrated in Figure 1(c): that of Figure 1(b) would be the result of amplifier properties which are absent from the definition given in Figure 1(a). How would we tell which of these two outputs is "correct"? We would investigate the response of a *real* amplifier (!): the *abstract logic* used to assemble the definition of Figure 1(a) is insufficient to satisfactorily integrate the different components into a self-consistent and globally correlated whole. *Abstract logic* is insufficient for unification.

2.2. Systems are unified by QM entanglement

So, if the *abstract logic* associated with Newtonian modeling is insufficient for system unification, *how* are systems unified? We must look for an inter-communicational logic which does support local-to-global-to-local correlation. Fortunately, we do not have to look very far – this is *precisely* the central nonlocal character of quantum mechanics.

The clearest example of this local-global-local coupling is given by quantum mechanical entanglement (QMe). If two Newtonian particles are separated by a large distance their inter-communication is restricted by the limited speed of light. Not so entangled quantum "particles". If two spin-correlated particles are produced together and then separated by a large distance, changing the spin of one of them results in an *immediate* change in the spin of the other, no matter what the distance is between the two. At a quantum mechanical level the differentiation of entities across the universe is correlated by QMe, as a force for unification. This resolves spatio-temporal discontinuities which could otherwise be created by the final merging of conflicting events whose occurrences are initially isolated from each other by communication-limiting relativity. Any doubts that QMe can apply to macro- rather than only micro-scale systems should be dispelled by the publication by Ghosh et al [4] of large scale QMe effects in magnetic systems.

We conclude that the central characteristic of *any* system is that its individual elements and scales are all unified in a single procedural-structure by quantum-mechanical entanglement. Where this is naturally available between the various parts of the system it can occur internally. Where it is *not* naturally available, for example in fragmented high level information processing systems, unification must be provided by other means.

2.3. High-level systems always contain life

It is maybe rather disconcerting to realize that a pictorial work of art does not itself "contain" the subject which we attribute to it. While this lack is reasonably evident in, for example, Picasso's painting "Nude Woman" in Figure 2(a), it is much less so in the photographic image of Figure 2(b), where we presuppose on the basis of common agreement that there *is a picture*. Both, however, share the same characteristic – the subject matter may well be present in the mind of the creator and in that of the beholder, but it is absent from the work of art itself, which has simply the nature of a communication channel. The clearest example of this nature is the transmission of a picture from studio to sitting room by television. The picture's image is focused on a pixilation device, the result transmitted as *independent* pixel details to the television, and then these are reassembled on the television screen *in the same spatial sequence as they were recorded*. The television "system" itself, however, could not care less whether the reassembly coincides with the original pixel order or not – it has no internal representation of interrelations between the pixels other than that imposed on it by the hardware manufacturer. A Boolean computer display has the same character: we see what we are *led to believe* is there, and not what *is* there.

Figure 2: (a) Picasso's "Nude Woman", (b) a giraffe.

Returning to our target of *abstract logic* information processing assemblies, what is it then that integrates their fragmented parts into a unified system? The nature of their isolated *abstract logic* does not permit this integration to take place directly between their individual components, but if they operate in a unified manner they must ultimately be somehow integrated by QMe. We conclude that it is the QMe of our own brains which provides this unification, by way of our integration into information-processing systems. The parallel processing required for perception of the subject of a painting is our own parallel QMe! High level information processing "systems" are only *systems* if *we* are a part of them. High-level systems always contain life.

3. WHERE DO GOALS COME FROM?

The most noticeable aspect of both James Albus's and David Fogel's descriptions of *intelligence* is that they both focus on the fulfillment of goals. Intelligence is both anticipative and intentional. In a dynamically evolving context it represents the capacity to correlate externally and internally projected versions of the future. Its application is recursive through learning, which embeds evolution into intelligence itself.

But where does this capacity come from? Are both its incidence and its recursivity unique to living organisms?

3.1. Both organisms and computers have instincts

Newly-born animals have instincts which enable them to survive and learn. These are built up from conception to birth as a pre-structuring of neural connections. A high degree of plasticity remains, however, enabling the animal not only to build on these instincts, but to replace them in many cases with environmentally-derived variants. Animals clearly demonstrate inborn instincts: the instinct to select energy-giving sugary and fatty foods; the instinct to learn to walk; the instinct to seek company; the instinct to group together for safety; ultimately the instinct to survive. These are observations of instinct at a holistic animal level, but we can find their precursors at more primitive levels. If the biological cell is the most primitive living organism, then the most primitive organism instinct is *cell mitosis*, whereby a single fertilized cell divides into two identical daughter cells, providing the means for assembling billion-cell organisms which demonstrate holistic instincts. It is important to note that although these high-level holistic instincts may be associated with *abstract logic*, in that they can be inconsequential at a late stage of organism evolution (e.g. a human baby's immediately post-natal attempts at walking), cell mitosis is a primitive *natural logic* instinct.

Computers display related properties which may be compared to instincts. The immediate post-power-up reaction of a computer is to load the basic input-output system (BIOS), which enables the computer to "survive" by adapting itself in a pre-programmed manner to its environment. This is clearly again holistic level *abstract logic* construction – and it is far from evident whether a comparison between it and animal instinct is laughable. But do computers possess primitive precursors to holistic "instincts"? Again, we need to look at the primitive structural and dynamic foundations of high-level operation. An obvious candidate for primitive computational "instinct" is that of electrons to follow the "rules" of solid state physics. This is again within the domain of *natural logic*. Both organisms and computers have *abstract logic* holistic and *natural logic* primitive "instincts".

3.2. Organisms replace instincts by goals

During the first few months of its life, a human baby spends endless hours violently kicking with its legs. Why? - it does not need them to walk. At least, not yet, but later on when it finally gets to its feet it will need strong leg muscles to support it. Instincts finally give way to goals. Primitive *natural logic* instincts give way to unreasoned possibly *abstract logic* goals, which in turn give way to *reasoned abstract logic* goals, which are subject to intentionality.

Ultimately, in the most neurologically-complex organisms, goals may be subjugated to motives. A baby's kicking provides the basis for walking, which clears the way to succeed in gymnastics or dancing, which may become a gateway for social contact and valorization of the individual. A nice example of way a social "motive" may override a goal is given by "the cockroach experiment". If a very small thirsty child is shown a three beakers of water, one clean, a second into which the child has seen a dead cockroach dipped, and a third still containing the cockroach, it will drink from all three. A few months older, it will drink from the first two, but not from the beaker which contains the cockroach. Again a few months older, it will now avoid drinking from the beaker into which it saw the cockroach dipped: the child has learned about contamination – the basic instinct to drink has been overridden by a fundamental social concept destined to reduce disease and promote survival.

The clearest example of replacement of instincts by goals is the recent upsurge in computer combat games. Here there is a direct replacement of the *natural logic* (i.e. consequential) fight for survival required of primitive mammals by the *abstract logic* (i.e. inconsequential) "fight for survival" against computer-generated enemies (it is interesting to note that one such game is publicized as "brutal combat for the thinking man"!).

Goals are created from evolving environmental information by the co-evolution of intelligence [1, 2], whether the information is from unintended uncontrolled occurrences or intentional social or instructive implantation. The major question of an entity which seeks to achieve or maintain states of its environment is "whose goals?"

3.3. Computers are fragmented, goals are not

Computers as we know them *do not* and *can not* contain *goals*! Their formal style of logic precludes any integration of their data, and even the individual bits which make up the representation of a single number are formally separate and devoid of meaning in the global context. Goals, however, if they are to be related in any consequential manner to detail of the environment within which they will operate, consist of numerous differentiable but intimately integrated components. Fragmented, purely *abstract logic* goals have little chance of being satisfied.

Computers as we know them are *not* integrated systems on their own, although they often appear to be so because we inadvertently include *ourselves* in their operation! The first criterion for a computer is that it must be capable of doing nothing - otherwise how would we know it is doing what we want it to do? Notice that this removes *all* autonomy from a computer. Situations where a computer appears autonomous are simply those times when the (formal) complication of their operation is too great for us to comprehend in detail and our (formal) simplified description of their operation is incomplete. Computers are fragmented, goals are not.

4. WHY ARE ORGANISMS DIFFERENT FROM COMPUTERS?

Let us look at biologically-derived intelligence as a prototype for artificial systems. Biological information processing *is* integrated: note the singularity of our individual consciousness. However, now we apparently find a contradiction. Our bodies clearly operate under the constraints of *natural logic*, but our minds do not appear to do so. Survival demands that we relate to our surroundings through simplified representations, as a way of reducing information processing time, and we consciously do so using *abstract logic*. But how does this come about?

4.1. Organism survival depends on abstraction

John LeDoux [5] has described how the mammal brain gains a temporal advantage (through "fear-learning") by processing high-risk information very rapidly via the amygdala rather than accurately via the cortex. Any successfully surviving entity needs to have available a wide range of differently scaled pre-computed "scenarios" [6] to enable it to react suitably to stimuli whose requirements only become evident within a restricted window of response. The necessary trade-off is between accuracy of response and speed of decision. These scenarios are *abstracted* forms of incoming environmental information, ranging from primitive models to extended representation.

LeDoux [5] gives as an example the way in which we are able to jump to escape a curved brown shape on the forest floor without waiting for the cortex to tell us whether it is a snake or not. Similarly, it makes no sense to examine the details of a car which is bearing down on us – better to first jump out of the way. But it is important, also, to be able to recognize if the driver is a family member, who is coming to pick us up and take us home! Recognition of both of these situations depends on *abstraction* from incoming information. It is important to note that the greater the information reduction in an *abstraction*, the greater its autonomy from other more elaborate representations of the same information: *abstract logic* can be generated by *abstraction* from *natural logic*.

4.2. Organisms link together natural and abstract logics

Living organisms structure themselves hierarchically, to circumvent organizational and cohesive difficulties which would otherwise cause their demise[1]. Probably the prime rationale is that by doing so they can construct high-level forms which are temporally quasi-independent of the low-level elements from which they are derived. Whereas the (high-level) output of a digital computer must wait for the

[1] There are some exceptions to this rule, notably among the fungi, where rather specialized criteria for survival pertain.

termination of every single one of its logic gate operations[2], an organism is able to react at its highest level without being impeded by lower-level slowly-evolving processes. The extension of neural glial-cell inter-communication times to some hundreds of milliseconds, for example, does not inhibit the brain's more rapid reaction to some stimuli. One major result of hierarchy is the partial autonomy which is associated with this partial *en*closure of hierarchical levels. John Collier [7], for example, has suggested that the brain cedes bio-support autonomy to the body, in return for extended autonomy of information processing. The multiple hierarchical levels we associate with an organism are a serial progressive *abstraction* from *natural logic* at the lowest levels to *abstract logic* at the highest ones.

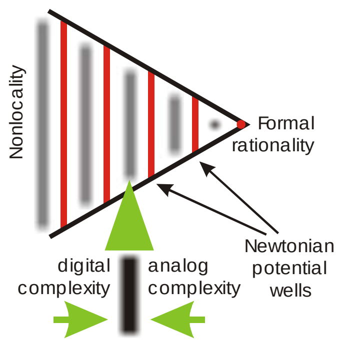

Figure 3: The general form of an organism hierarchy.

We have previously published extensive details of specific [6] and general [8] models of hierarchies which correspond to this formula. Figure 3 illustrates the general outline of a hierarchical organism. Differently scaled perceptional levels (indicated by the vertical lines) in the form of Newtonian potential wells in a general phase space are coupled together by quantum mechanically-related inter-scalar regions where the necessary inter-level complexity is pooled. If a hierarchy of this type is permitted to self-stabilize by inter-level correlations across the entire assembly, then the inter-scalar regions correlate as well, into scalar forms of a rationality which is complementary to that of the Newtonian wells. Figure 4 illustrates this segregation into two quasi-independent rational hierarchies. The Newtonian set is reductive towards localization; the inter-scalar set is reductive towards nonlocalization: these are differently scaled forms of QMe. Here again, all the elements of a unified assembly are correlated by QMe, not only at a specific scalar level, but also between scalar levels.

4.3. Computers use isolated abstract logic

If both organism and computer start off from *natural logic* and develop towards the use of *abstract logic*, where is the difference between them? The developmental progression from *natural* to *abstract logic* in an organism is continuous. Although there are strong environmental and external directive effects, the development itself is wholly internal: an organism "does it itself". *Natural* and *abstract logics* are integrated within the organism, and ultimately the *abstract* can only be divorced from the *natural* by internal decision (quasi-autonomously or not).

A good example of the way organisms manipulate their internal linking between *natural* and *abstract logics* is given by the way we learn to play games which require the integration of various physical actions, such as golf. The process of learning a "golf swing" is initially very close to the *natural logic* level, where we concentrate on individual detail and techniques to evaluate their prospective part in ultimate success and to model them *abstractly* (i.e. "if I do 'this' and 'that', and then 'that', the result will be 'that'). However, this is not how 'good' competition golf is played! The required follow-up is extensive practice, until the detail becomes integrated into a complete action which can be performed "naturally" without attention to the detail: this entails *en*closure of the detail into a form where it can be consciously addressed in a more symbolic manner: the details have been returned to a more *natural logic* level where they are contextually *integrated* and can undergo inter-detail compromise. Unfortunately, the stress involved in competition makes us again focus attention on specific details, and the integration breaks down, leaving us to make *abstract logic* errors without any obvious reason. Stress-reducing drugs such as beta-blockers are consequently understandably unwelcome in competition golf, as they permit players to maintain the *natural logic* integration of detail and the inter-detail compromise which results in an 'automatic' good swing.

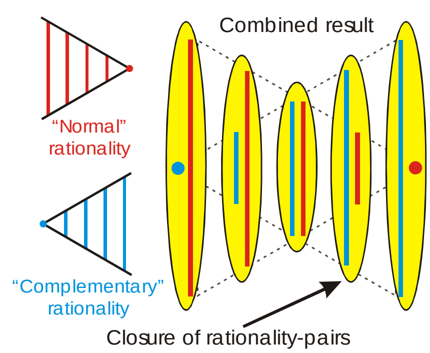

Figure 4: The separation of a organism hierarchy into complementary rationalities.

Within a computer there is *no* linking of *natural* and *abstract logics*: their segregation is the *first rule* of computer design, to take decision-making out of the computer's "hands" and keep it for ourselves.

It is important to remember, however, that a wholly *abstract* level of cognitive effort is far from useless. Nature has endowed us with the capacity to dream, not only as a

[2] The inclusion of an IF..THEN clause does not change this, as the consequently irrelevant part of the program does not then "exist".

way of re-correlating an evolving global neural database [9], but as a way of creating and visualizing feasible scenarios without experiencing unpleasant consequences. Creativity depends on the free run of imagination. By its very nature, inventiveness initially operates on the borderline of "the consequent", until such time as its machinations compose a formulation which can be linked into the natural logic of our environment as "a practical proposition".

5. How can we make Virtual Organisms?

So, if we are to build intelligent goal-driven systems, how should we do it? The current approach of taking a system which by its very nature includes *us* in its operation seems rather strange. Artificial life does not exist: it is nothing other than a functional simulation of life, projected into an *abstractly*-operating embodiment by us, the godlike designers. Can we see how we should initially proceed in creating virtual (v)organisms?

5.1. Natural and abstract logics must be linked

In common with natural biological organisms, a virtual (v)organism must implement integrative linking between the *natural logic* of its primitive constituent elements and the high-level *abstract logic* domain within which we wish to embed (our) goals (for it). It is difficult to see any path towards this requirement which does not pass through the establishment of a birational self-correlating hierarchy of the kind we describe above, where system unification is assured by QMe. A virtual (v)organism will need both grounding in *natural logic* and the application of that grounding to its *abstract logic* operation. As we pointed out earlier, goals are created from evolving environmental information. The associated co-evolution of intelligence within a (v)organism will entail not only its adaptation to a changing environment, but also the co-evolution of its goals. The question of "whose goals?" becomes a very relevant one. Do we *really* want to create autonomous (v)organisms, or would we prefer that their goals are in any case related to our own desires? Operational goals must be derivable from instincts (or whatever we wish to call the "start-up" criteria for such a system). The goals which we initially implant into a (v)organism may be the seeds of its downfall if we do not permit the co-evolution of environment, intelligence and goals. It will clearly be necessary to develop a progressive approach to the establishment of a (v)organism's mode of operation and its goal selection, if we are to remain in any sense associated with its intelligence.

5.2. Autonomous (v)organisms need parents

Nature has developed the quasi-integrated role of *parent* to its child. Maybe this is the position we should seek if we wish to expand our capabilities through developing quasi-autonomous (v)organisms but still retain control of them. The *biological* relationship is not one of subordination, or of segregation, but one of partnership. Child learns from parent, parent (hopefully) learns from child. It is a relationship of co-evolution, where one (the parent) can provide numerous learned "quick and easy fixes" to environmental problems, but where the other (the child), by dint of its lack of pre-formulated scenarios, is freer to generate new and original approaches, and in doing so to modify its approach to the environment, and the environment itself, more radically.

As is usual with human parent-child relationships, the designations of "parent" and "child" cannot necessarily be automatically attributed to one or the other of the partners on the basis of superficial appearance!

6. A Final Note

In the light of the preceding discussion, it is interesting to note that it is required of this article that it be headed by an *abstract*. We leave it to the reader to decide whether he or she would agree with the authors that the *abstracted* content corresponds sufficiently well to the *natural* content of the article itself, and the extent to which its accuracy depends on the context within which it is viewed by the reader.

References

[1] Albus, J. S. "Features of Intelligence Required by Unmanned Ground Vehicles," available on the NIST publications list as http://www.isd.cme.nist.gov/documents/albus/Features_of_Intelligence.pdf .

[2] Fogel, D. B. "Evolving Solutions that are Competitive with Humans," in Performance Metrics for Intelligent Systems Workshop, Gaithersburg: NIST, 2002.

[3] Edmonds, B. "What is Complexity? The Philosophy of Complexity per se with Application to Some Examples in Evolution," in Heylighen, F., Bollen, J. and Riegler, A., eds., *The Evolution of Complexity: the Violet Book of 'Einstein meets Magritte' (V.8)*. Dordrecht: Kluwer Academic, 1996.

[4] Ghosh, S., Rosenbaum, T. F., Aeppli, G. and Coppersmith, S. N. "Entangled Quantum State of Magnetic Dipoles," *Nature*, vol. 425, pp. 48-51, 2003.

[5] LeDoux, J. E. "Brain Mechanisms of Emotion and Emotional Learning," *Neurobiology*, vol. 2, pp. 191-197, 1992.

[6] Langloh, N., Cottam, R., Vounckx, R. and Cornelis, J. "Towards Distributed Statistical Processing: a Query and Reflection Interaction Using Magic: Mathematical Algorithms Generating Interdependent Confidences," in Smith, S. D. and Neale, R. F., eds., *ESPRIT Basic Research Series, Optical Information Technology*. Berlin: Springer-Verlag, 1993, pp. 303-319.

[7] Collier, J. D. "Autonomy in Anticipatory Systems: Significance for Functionality, Intentionality and Meaning," in *1998 Computing Anticipatory Systems: CASYS'98 - 2nd International Conference, AIP Conference Proceedings.*

Woodbury, New York: American Institute of Physics, 1999, vol. 465, pp. 75-81.

[8] Cottam, R., Ranson, W. and Vounckx, R. "A Diffuse Biosemiotic Model for Cell-to-Tissue Computational Closure," *BioSystems*, vol. 55, pp. 159-171, 2000.

[9] Cottam, R., Ranson, W. and Vounckx, R. “Localization and Nonlocality in Computation,” in Holcombe, M. and Paton, R. C., eds., *Information Processing in Cells and Tissues*. New York: Plenum Press, 1997, pp. 197-202.